

# Peer-to-Peer Storage Systems: Crowdsourcing the Storage Cloud

Speaker:

**Anwitaman Datta**

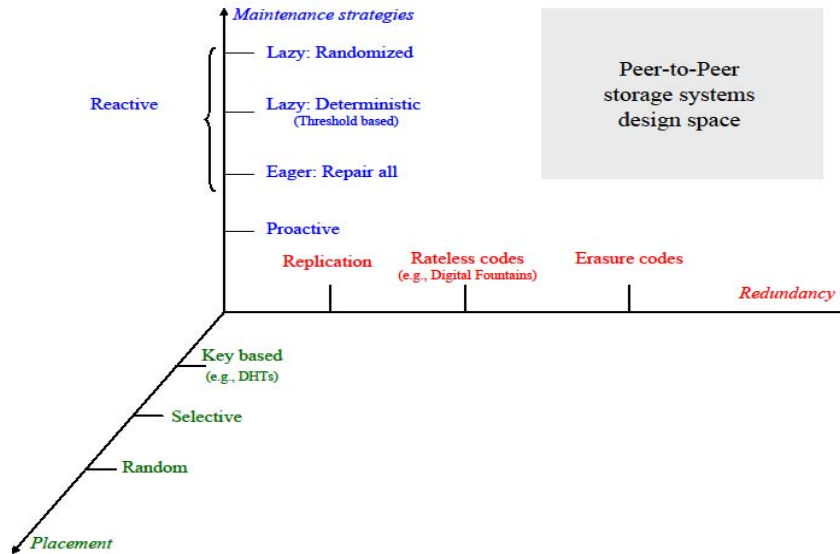
**Nanyang Technological University, Singapore**

## Abstract

This decade has witnessed vigorous research on peer-to-peer research – encompassing numerous areas including structured overlays, file sharing and content distribution networks and storage systems to name a few. Numerous books, tutorials and surveys on peer-to-peer systems exist, but none exclusively focus on storage systems, nor do they exhaustively explore the design space. This tutorial will summarize the fundamental design issues related to peer-to-peer storage systems, pertaining to the reliability aspect of system vis-a-vis availability and durability. Efficient data indexing and searching as well as security issues are also important aspects of storage systems, but will be out of the scope of this tutorial, but brief references will be provided.

Peer-to-peer (P2P) storage is a paradigm to leverage the combined storage capacity of a network of storage devices (peers) contributed typically by autonomous end-users as a common pool of storage space to store and share content, and is designed to provide persistence and availability of the stored content despite unreliability of the individual autonomous peers in a decentralized environment.

Crucial to peer-to-peer storage systems research are some of the following questions: What kind of redundancy is most efficient from various perspectives including storage overhead as well as access and maintenance costs and latency, and implementation complexity? What minimal redundancy is necessary to meet a desired level of resilience? Which maintenance strategy to apply to maintain the necessary redundancy? Which peers to store the redundant blocks in, based possibly on issues like reliability of individual peers, locality and load? The tutorial will provide a survey of related literature exploring these aspects of a peer-to-peer storage system design (the design-space is summarized in the figure below). In doing so, we will review a whole spectrum of literature, including analytical models and simulation based studies, as well as algorithmic and systems implementation issues and various applications such as archival and back-up, file systems, etc – and the specific requirements and flexibilities of such systems, which need to be understood and exploited for reliable, efficient and scalable systems designs. A more precise list of topics to be covered is outlined below.



## Tutorial Outline

### 1. The P2P Paradigm & background information

History and philosophy of the P2P paradigm

P2P paradigm from the perspective of distributed systems

Distributed index structures and key based routing principles (PRR, DHTs, etc).

### 2. P2P storage systems: Basic concepts & design space

Kind of redundancy – replication, coding (traditional erasures, rateless, network)

Redundancy maintenance (and garbage collection) mechanisms – reactive, lazy, proactive, ...

Placement strategies

### 3. Analytical models to understand storage systems behaviour under churn

Static resilience – Probabilistic models, dynamic equilibrium studies – Markov models, Dynamical analysis – Differential rate equations

### 4. Empirical studies of storage systems under churn

Effect of redundancy strategies, maintenance strategies and some results on placement strategies

### 5. Systems & applications

OceanStore/POND, TotalRecall, Glacier, Antiquity, FS2You, may be some others ...  
ePost, LOCKSS, CFS, PeerSoN

### 6. Open issues and research opportunities

## **About the Speaker**

Anwitaman Datta did his PhD at EPFL Lausanne before moving to NTU Singapore in 2006, where he is currently an Assistant Professor. He is interested in large-scale networked distributed information systems and social collaboration networks, self-organization and algorithmic issues of these systems and networks and their scalability, resilience, security and performance. He won the best paper award at ICDCS 2007 and is one of the recipients of HP Labs Innovation Research Program award 2008. He teaches a course on Advanced Topics in Distributed Systems in NTU, which covers most of the topics mentioned in the tutorial, and has published more than 30 papers including several on relevant topics for the tutorial, as well as other topics related to peer-to-peer as well as self-organizing systems, as well as some survey style book chapters relevant to subtopics of the tutorial, as listed below. He also serves as a program co-chair of P2P 2009.